Using BGP to attack the Internet

(and some IPv6)

Tarko Tikan
tarko *at* estpak.ee
TTÜ, 09.11.2010

This is internet

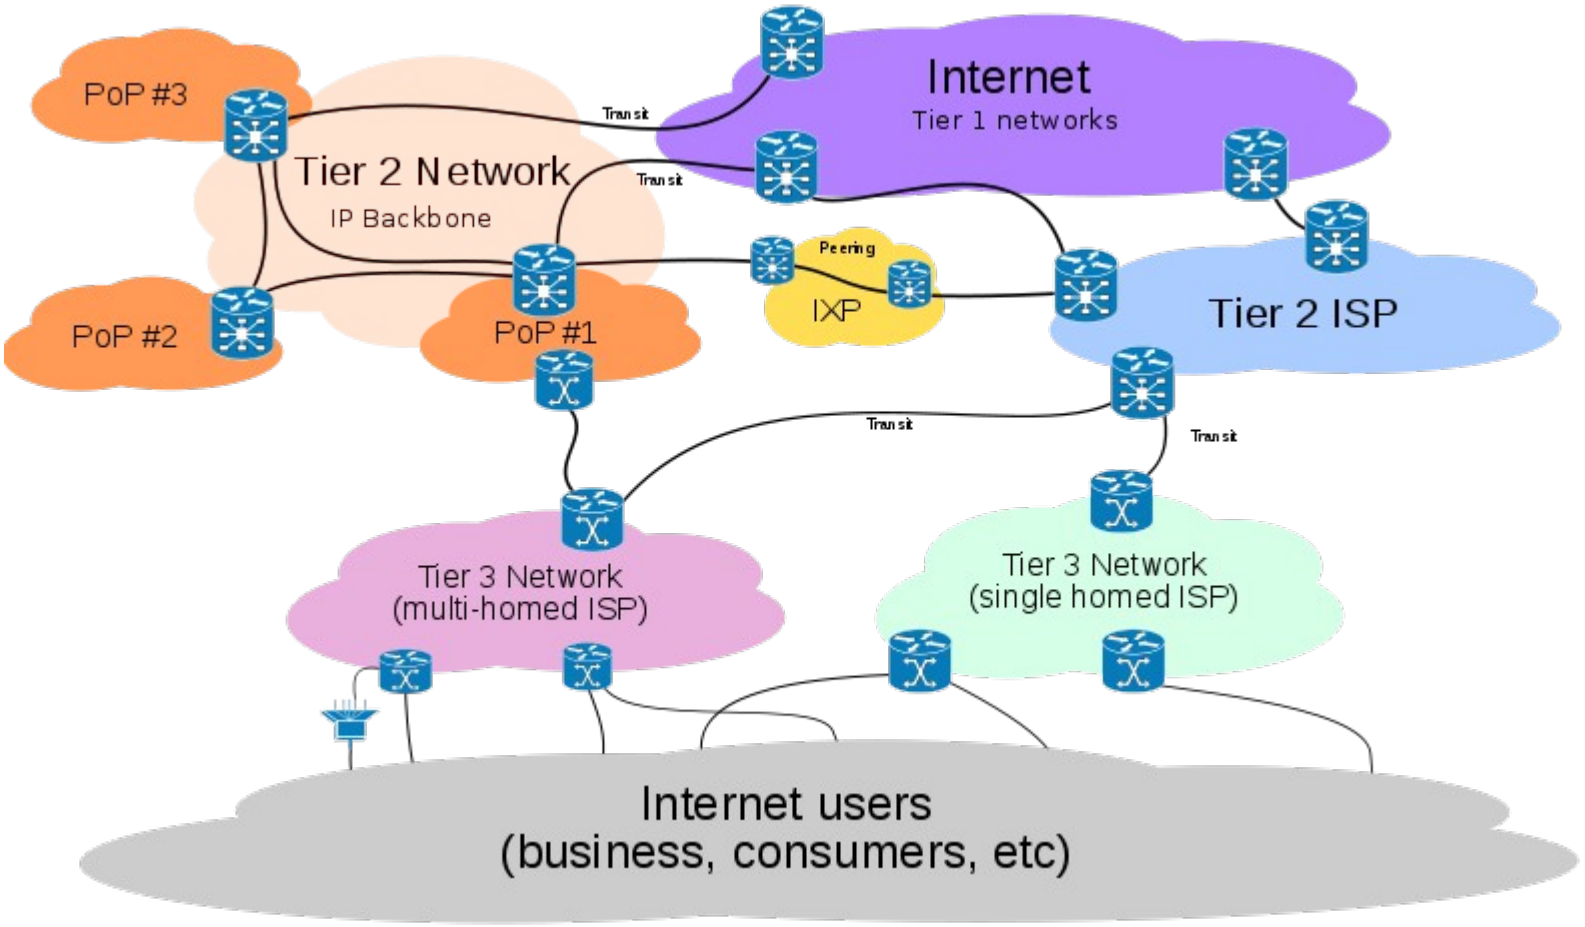

http://en.wikipedia.org/wiki/File:Internet_Connectivity_Distribution_%26_Core.svg

ISP

- Has:
  - AS number(s)
  - IP addresses (aka routes aka prefixes)
  - At least 1 transit connection
  - Possibly lot of peers and transits
  - Customers and/or content
    - Perhaps downstream customers who they speak BGP with

AS & IPs

- Needed for BGP (reverse is also true)
- Allocated by IANA → RIR → LIR
  - Example: IANA → RIPE (EU) → Elion
- Registered in RIR database (whois)
- Routing information also registered in RIR database (if registered)
- Same for contact information
- No country borders (myth of .ee)

Transit

- “Gateway” to full internet, ~338000 routes
- Paid by customer
- Usually filtered by upstream (unless customer is very big or upstream is stupid)
- “Tier 1” = having no transits

Peering

- Interconnection between 2 ISPs
- Zero cost (in reality: 50/50)
- Preferred over transit (cost, distance)
- Routes: peering partner's + it's customers
  - Filtering hard, one side can't demand, lots of work
- No up/downstream
  - Both sides want better deal
  - Filtering becomes tricky
  - Leads to peering wars, partitioning

Peering

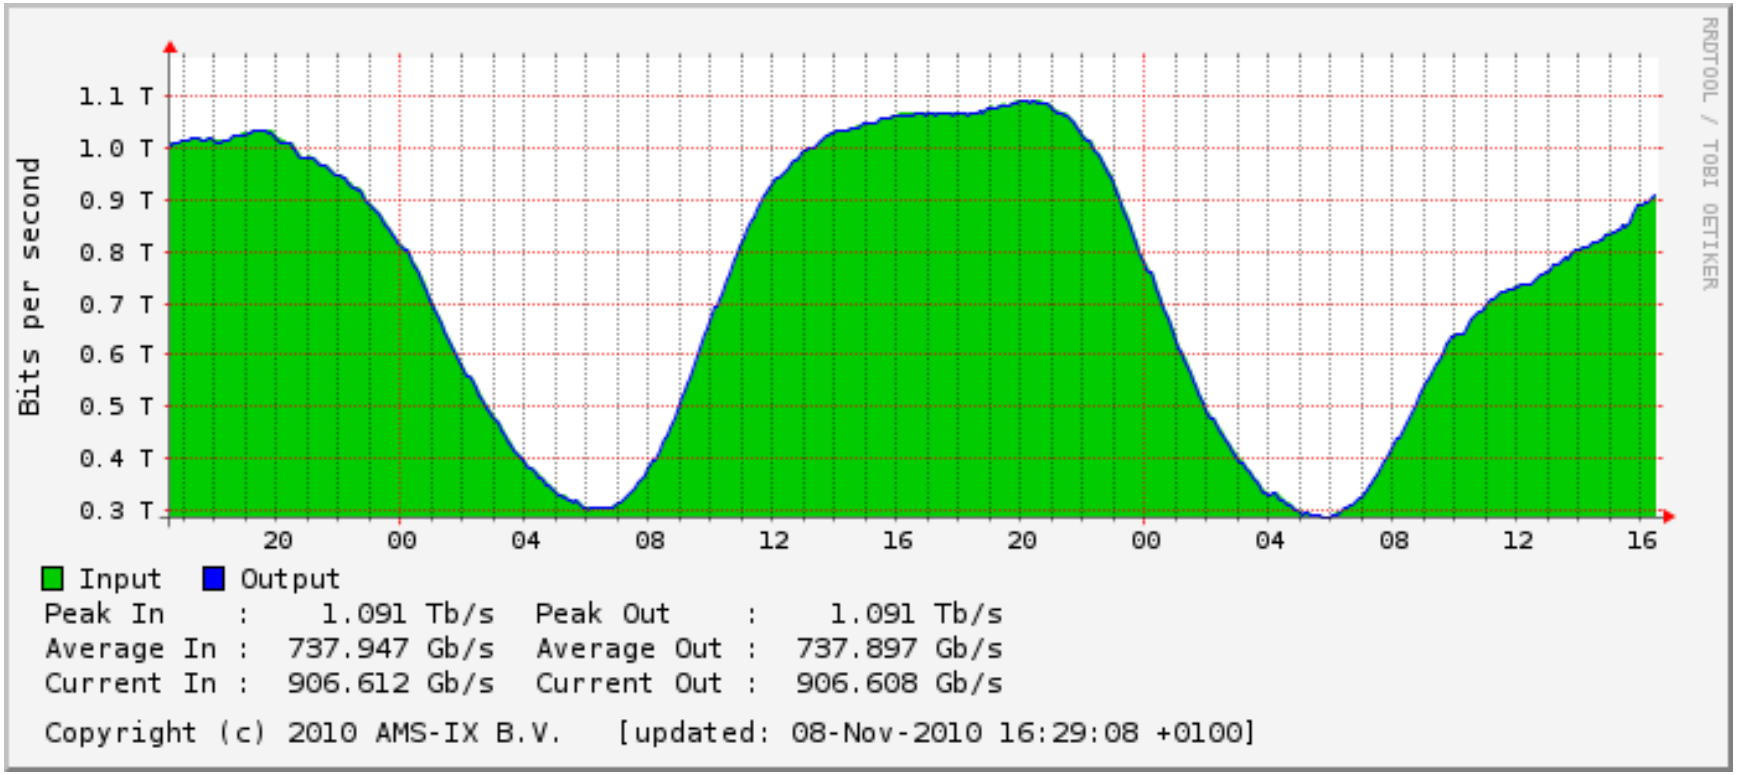

Also check: http://www.renesys.com/tech/presentations/ for analysis on tier1/(de)peering etc.

BGP simplified

- BGP is a routing protocol running over transit/peering/customer links between 2 parties (ASes)
- Used to advertise routes (prefixes) to neighbor
- Using TCP session over IP
- eBGP and iBGP (focus on eBGP)
- Also used on smaller scale:
  - BGP2server
  - VoIP etc.

What to defend against

- I don't want other end to kill my router
- I don't want other end to send me bogus routing information
  - Pollutes my view
  - Leads to reachability issues
  - Other nasty stuff I'll speak about
- Myself – I don't want to “kill the Internet”
  - Easier than you might think :)

Flood router

- Either number of routes or routing updates
- Why is this a problem? Router has limited memory for routing table, limited FIB size, limited CPU resources to process routing updates
- Can very easily be accidental – route leaking, number of incidents over years
- Mitigation – maxprefix
- http://www.renesys.com/tech/presentations/pdf/renesys-nanog34.pdf

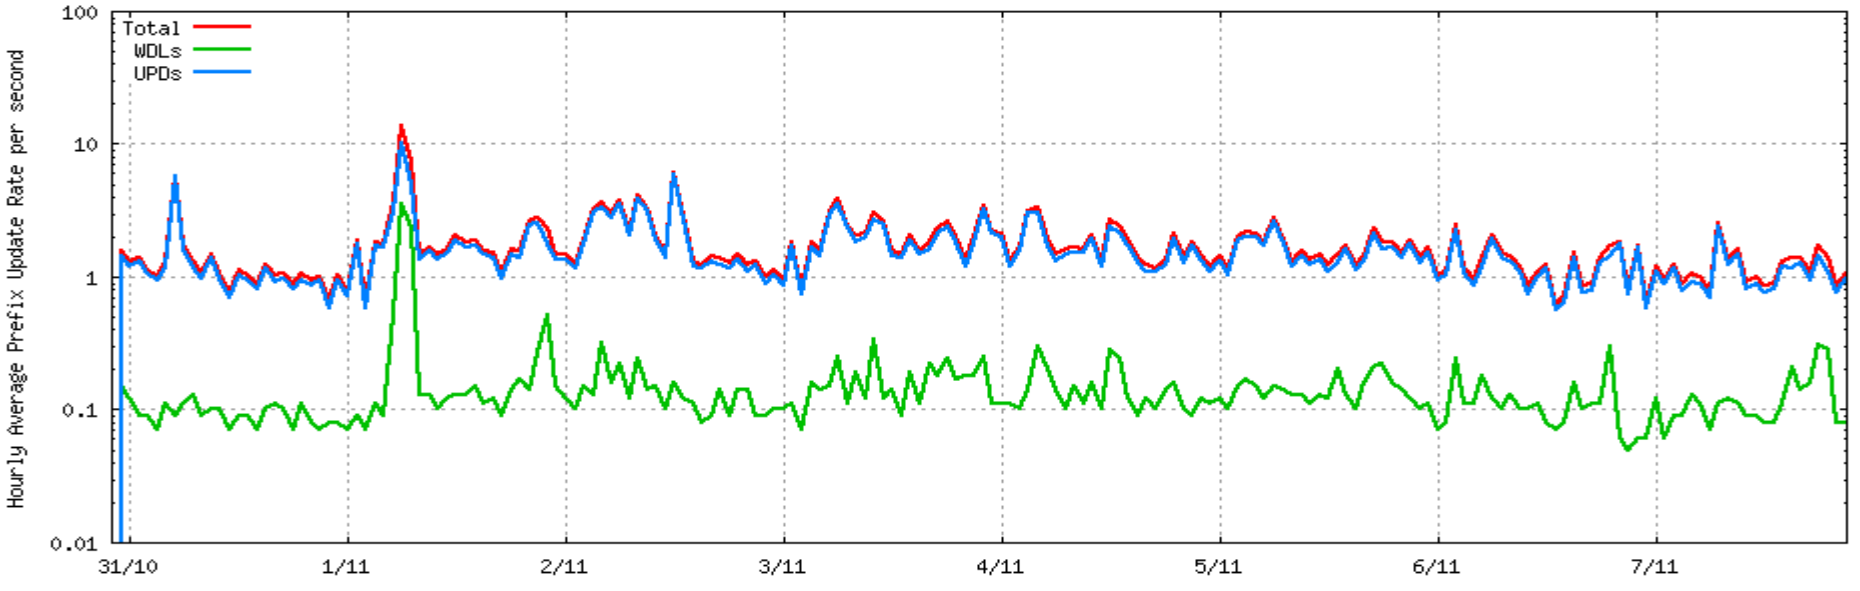

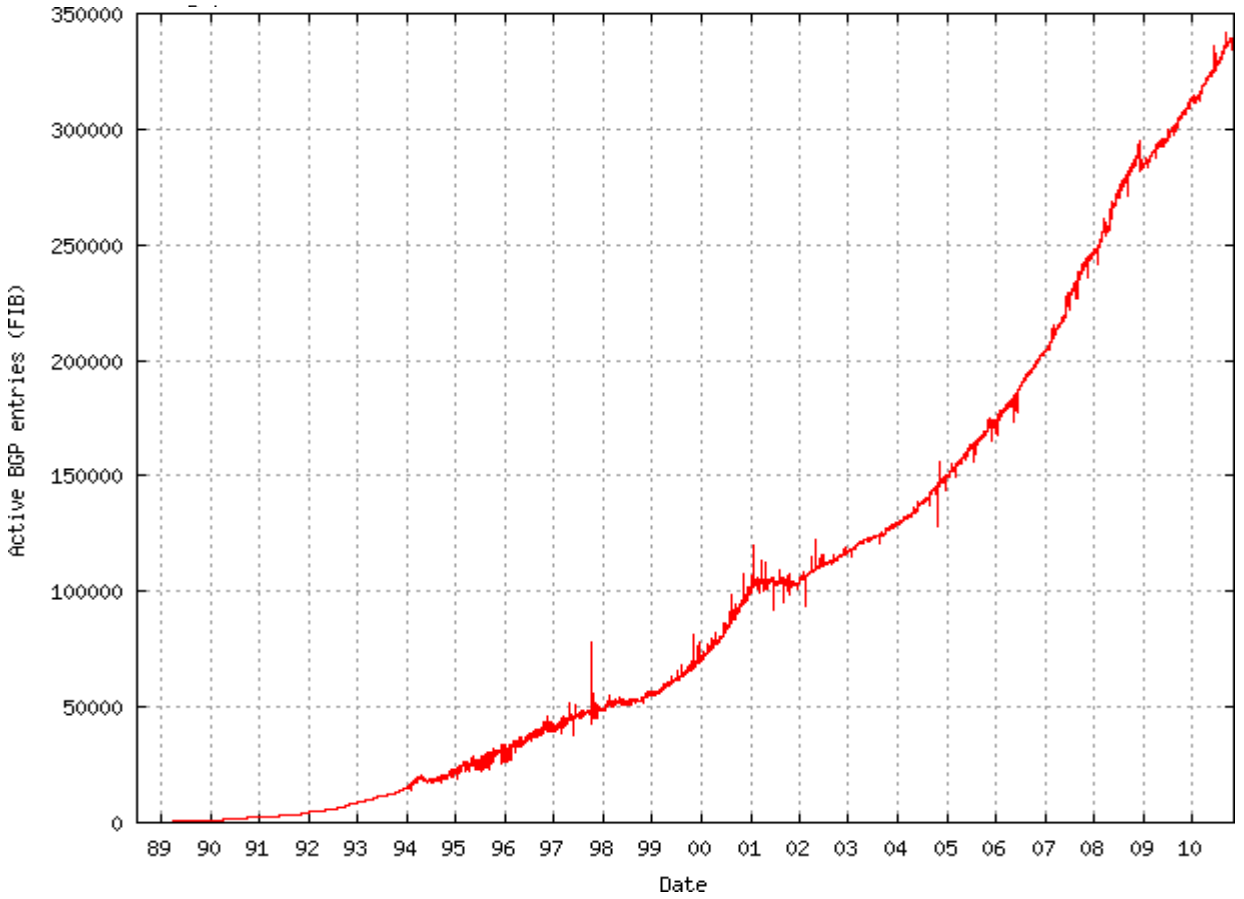

http://bgpupdates.potaroo.net
/instability/bgpupd.html

http://www.cidr-report.org

Attack TCP session

- As BGP is running on top of TCP session, all attacks against TCP are valid
- Guess src/dst ports and SEQ number and send TCP RST
  - Leads to TCP session reset
- Mitigations:
  - But no one can spoof source (Ha! Ha!)
  - Use TCP md5 based session authentication
    - Key change problematic
    - If under attack, makes situation worse
  - TTL security

Hijack address space

- Why this works?
  - More specific routes have preference
  - If announced to peer, routes get preference over his transits
  - /24 problem
  - Because how BGP works (announce everything by-default), very easy to accidently leak
  - Luckily, with routes, you also get traffic which overloads your network and BGP sessions drop :)
- Not theoretical, happens every day (mostly on small scale)
  - Youtube case best known
  - http://www.renesys.com/tech/presentations/pdf/menog3-youtube.pdf
- Has good uses also
  - Quick filtering inside network, DDoS protection
  - If you accidently announce these routes outside your network, you have a problem (Youtube case)

Hijack address space

- Can't you just validate routes and filter others?
  - How to figure out which routes are valid? (IRR trust problems)
  - Attacker can hijack database as well
- Even if you know which routes are valid:
  - Long prefix lists create interesting effects (juniper commits, cisco nvram size) and slow down convergence
  - Prefix list is not ideal, no path or origin validation
  - Proposed solutions include lots of “certificates, root, PKI, trust anchor” buzzwords
    - There is no root trust
    - Your router can't depend on external system

Hijack unallocated address space

- How is that different from previous?
- Mainly used to:
  - Send spam
  - Hack someone leaving almost no traces
  - Both cases → short living routes
- Route collectors deployed (wayback machine):
  - http://www.ripe.net/projects/ris
  - http://www.routeviews.org/
- But I only use the space inside my own network?
  - Lots of people do that
  - Reachability problems in the future (as of october, 12 /8 left)
  - 1.0.0.0/8 allocated to APNIC in january 2010
    - Receives ~200Mbps of traffic
    - 1.1.1.1 is worst :)

Send broken BGP message

- Lot of problems with BGP attributes and message parsing in general
- This won't kill just one router
  - If attribute parsing problem, update is propagated
  - 1 AS (vendor), multiple ASes, whole internet :)
- Recent cases:
  - Mikrotik & Cisco – software & operator error
  - Amsix & CRS-1 (limited scope – peering) – result of experiment
    - http://labs.ripe.net/Members/erik/ripe-ncc-and-duke-university-bgp-experiment

Why use any of these attacks?

- DoS someone
  - Yes, DoS, all you need is one BGP session and clueless operator
- Steal service
  - DNS is best – steal DNS, redirect all records to malware distribution site
  - Even root DNSes? Anycasting to fight the problem
- Intercept traffic (MITM)
  - Not that easy but doable
  - http://nanog.org/meetings/nanog44/presentations/Tuesday/Kapela_steal_internet_N44.pdf
- Luckily 99% operator/vendor errors
  - But spam sending is very real

Summary

- Can't introduce new protocol
- Slow adoption
  - No single organization rules the Internet
  - Lots of old hardware/software used
  - Most operators clueless and afraid of change
  - Router vendors default configs are no good
- It's only as strong as its weakest link

IPv6

96 more bits :)

Why?

- Almost out of IPv4 addresses
  - 12 /8 remaining

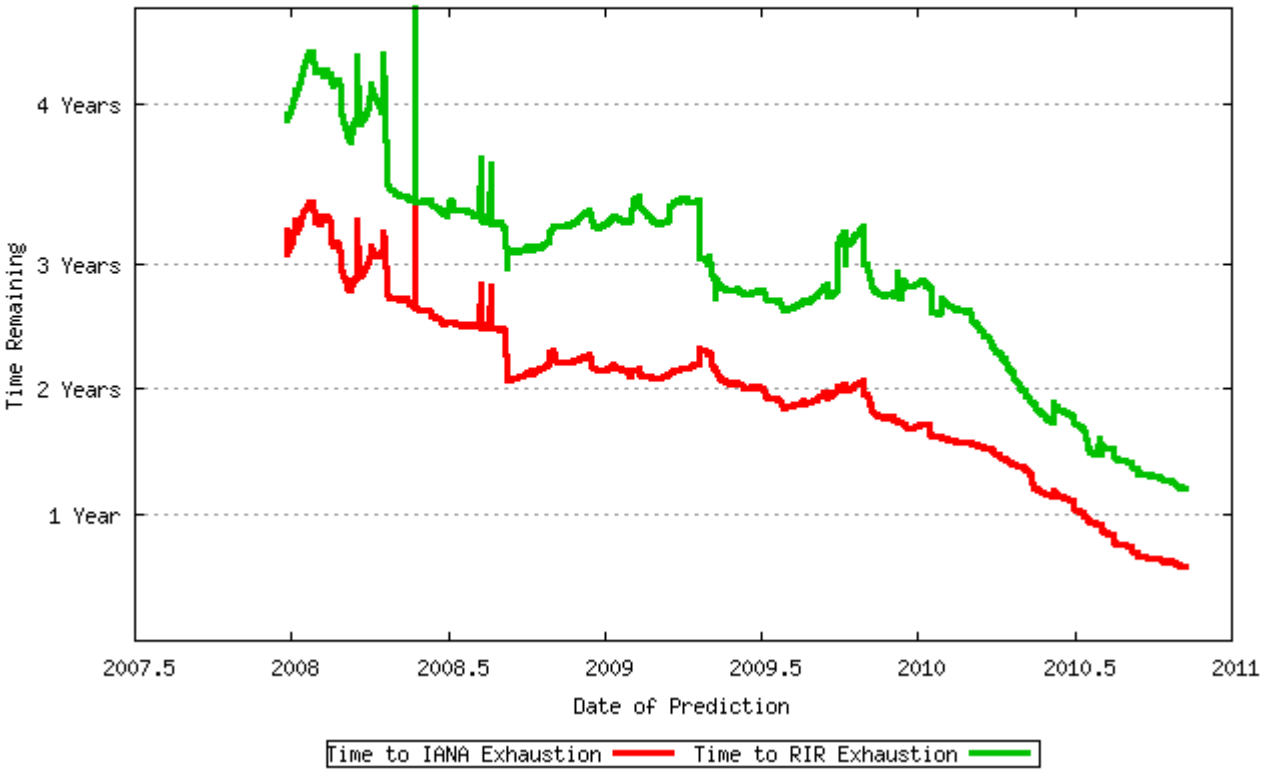

http://www.potaroo.net/tools/ipv4/

Doing things ~~right~~ this time

- Wanted to do things right this time but as always...
- Lots of new features no one is using and things that are already deprecated
  - ip6.int > ip6.arpa
  - /35 > /32
  - Routing header type 0 - the good old source routing
- In the end: ipv6 is not mass deployed and you already need a lot to be secure

No less secure than IPv4 :)

- But not more secure either
- ND and RA, ND much more complicated than ARP
- Ethernet L2 attacks still the same but old L2 securing measures useless
- NAT or no NAT - will stay, renumbering & multihoming
- Address space scanning complicated
- RA don't support DNS, DHCPv6 is yet again new protocol
  - Yes, I know, RA has extensions but no OS supports it today

But they told “security” is mandatory

- IPv6 has IPsec!
  - IPv4 has IPsec too
  - IPv6 mandates IPsec but doesn't mean anyone is using it
  - Still as complicated as for IPv4
  - Number of places you could use IPsec to workaround problems
    - Key management issues

Going wrong with IPv6

- RA giving addresses to all machines in subnet
- No IPv6 firewall in place
- Autotunnels
- Services bound to :::xxxx listen for IPv6 as well
- IPv6 transition using NAT
  - This will happen
  - Address sharing will bring all new issues
- Everything said about BGP is still true

You have reached the very last page of the ~~Internet~~ presentation

In case of questions (please allow few days for a reply):

tarko *at* estpak.ee